CS 289 Knowledge representation and reasoning
Fall 1995 Stuart Russell

Questions on readings

Week 1 (8/28, 8/30): Artificial intelligence and knowledge representation

Russell & Norvig, Exercises 2.1, 2.3.

Week 2 (9/6): Propositional logic

Russell & Norvig, Exercises 6.1, 6.3a–d, 6.5, 6.7.

Extra credit: derive an analytical expression for the "critical ratio" of clauses to variables for which the probability of satisfiability is 0.5.

Week 3 (9/11, 9/13): First-order logic

Russell & Norvig, Exercises 7.2, 7.11.

Why does Levesque pinpoint intractability as the central issue in KR, and what possible solutions does he discuss?

Week 4 (9/18, 9/20): Representing the commonsense world

Russell & Norvig, Exercises 8.2, 8.6.

What are the pros and cons of Hayes's representation of liquids via the place occupied? What sorts of facts cannot be expressed in this representation?

Extra credit: find one or more inconsistencies in Hayes's axioms.

Week 5 (9/25, 9/27): Inference algorithms

Russell & Norvig, Exercises 9.7, 9.8.

How can goal-directedness (as in backward-chaining) be obtained in resolution-based systems?

Given a sentence α and a theory B, how might one construct all theories H such that $B \wedge H \models \alpha$? Can this be done in general?

Briefly describe the modifications made to Prolog to yield PTTP.

Week 6 (10/2, 10/4): Complexity of logical inference

Are expressiveness and tractability monotonically related?

What happens to problems such as TSP in a polynomial-time logic?

Give common-sense examples of concepts in $\mathcal{PL}_1$ and $\mathcal{PL}_2$ (Donini et al.), making full use of the expressive power of those languages.

Define the term "approximation" à la Selman and Kautz.

Can offline computation and compilation have any bearing on the impact of intractability on expressive KR systems?

Week 7 (10/9, 10/11): Planning and plan execution

Russell & Norvig, Exercise 11.2.

What precisely are the differences between Blum and Furst's Planning Graphs and the partially ordered plan structures created by POP? How do these differences contribute to the efficiency of the algorithm? Comment on the experimental methodology of the paper.

How are action failures detected and handled in Sage? What happens when the domain failure is fixed? Suggest a way in which appropriate monitoring actions could be incorporated by the general planning process.

Week 8 (10/16, 10/18): Alternative agent architectures

What are Brooks's main arguments? What is wrong with them?

Where do Pengi's environment and agent design fit into Russell & Norvig's classifications?

Characterize the Botworld according the Russell & Norvig classification scheme. What kind of agent is a TR agent? Discuss the relationship between a TR plan and an annotated plan for execution monitoring, as described in Russell & Norvig Chapter 13.

Week 9 (10/23, 10/25): Software agents

Summarize in PAGE form the design problem for software agents, and explain how the various proposals in the five papers contribute towards a solution.

Describe the form, content, and role of the knowledge represented in each of the agents described.

Week 10 (10/30, 11/1): Propositional and first-order probability

What are the operational differences between a logical theory that's approximately correct and a probabilistic theory?

Russell & Norvig, Exercises 14.5, 14.8.

Prove from first principles the rule of conditioning:

$$P(A|B) = \sum_c P(A|B, C{=}c)P(C{=}c|B)$$

Do Bacchus, Halpern, and Levesque discuss actions with uncertain outcomes? Explain, with an example, how these can be handled.

Week 11 (11/6, 11/8): Probabilistic networks

Russell & Norvig, Exercises 15.1, 15.3.

Loop cutset conditioning creates a set of "copies" of the instantiated network. Explain how solving these copies might lead to repeated work, and suggest a way of avoiding this.

Explain how subnetwork extraction is done in CPCS, and what it means for this to be a sound process.

Week 12 (11/13, 11/15): Probabilistic networks contd.

Discuss prospects for parallelization of the various algorithmic approaches to belief network evaluation.

Describe a network for which the algorithm of Delcher *et al.* would result in a speedup of 1000.

Augment the car insurance network from Russell & Norvig Exercise 15.5, shown below, with appropriate QPN annotations. Comment on any qualitative synergies you observe.

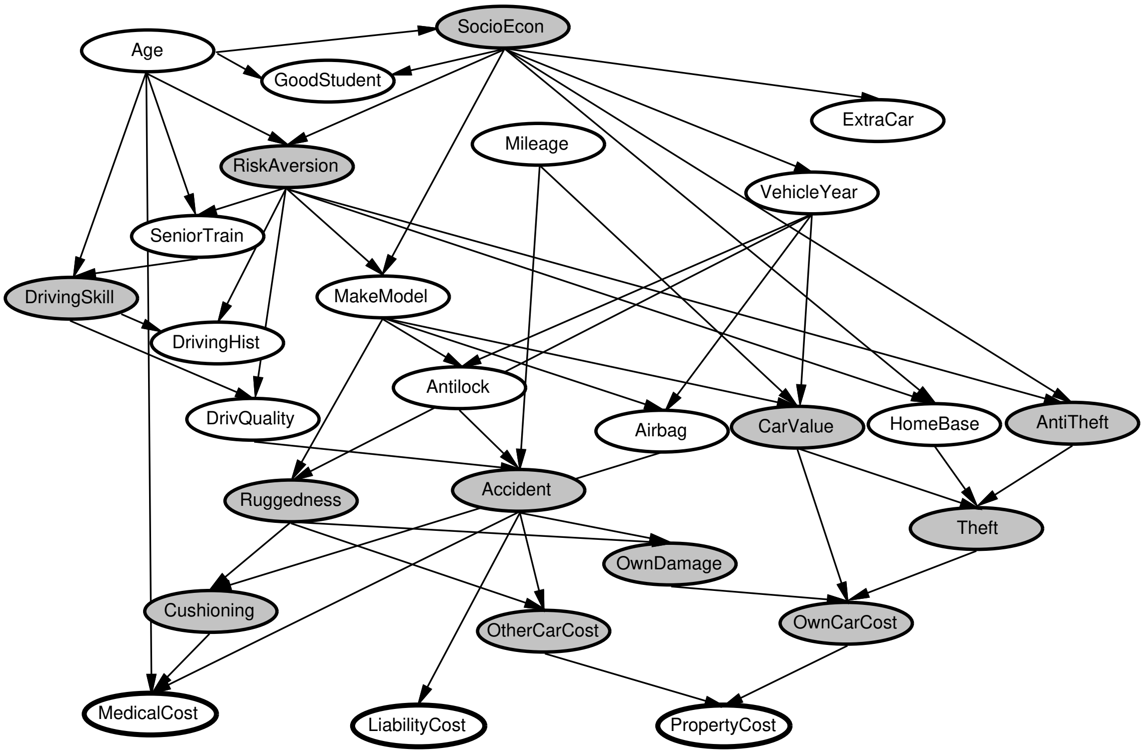

Week 13 (11/20, 11/22): Decision theory and sequential decisions

Russell & Norvig, Exercises 16.3, 16.11.

Suppose one applies the approach of Dean *et al.* to a stochastic version of the 8-puzzle. How might one incorporate heuristic information, e.g., the Manhattan distance heuristic, into the algorithm?

Discuss how you would go about constructing a utility function for an automated taxi.

Week 14 (11/27, 11/29): Partially observable environments

Russell & Norvig, Exercise 17.6.

A belief state is a probability distribution over world states. Why is it incorrect to say that the utility of being in a belief state is the expectation, according to the probability distribution, of the utilities of the world states?

How to DPNs help in designing agents for partially observable environments? How is the belief state represented in the DPN?

Week 15 (12/4, 12/6): Decision-theoretic planning

What are the relative advantages and disadvantages of policies vs. conditional plans?

What does clobbering mean in a probabilistic context?